

Please amend the specification as follows:

Please replace the paragraph beginning at page 8, line 26, with the following:

---

A<sup>c</sup> --In still another preferred exemplary embodiment, the local database of record at each service processing node is maintained by at least one data management component. The data management component makes managed objects and other data available to service processors at the service processing node. In particular, the data management component serves as a persistent store for managed objects, operating data such as adjacent node connection information, and other data created during service processing. Data management component also controls the replication of managed objects within various service logic execution environments (SLEE's) at the service processing node. By virtue of maintaining a local storage, the data management component affords fast restoration of normal operations after, or in response to, a network equipment failure, maintenance switchover, or other similar circumstances. For example, if a service processing hardware element suddenly fails, the service processing context can be replicated in another processor and the service resumed without interruption. Having the data and managed objects stored locally at each service processing node provides for fast fail-over and recovery without the delays associated with retrieving managed objects from the network service administrator component.--

---

Please replace the paragraph beginning at page 18, line 5, with the following:

---

A<sup>2</sup> --As shown conceptually in Figure 3(a), the Service Administration component 500 is a component that performs all of the functions needed to manage, store, and distribute all services and service data used by service processing nodes and to configure both the hardware and

A software components throughout the network control system depicted in Figure 1. Generally, as show in Figure 3(a), the SA component 500 is responsible for: cataloguing and storing the created objects and data from MOCE (Service Creation) 228; returning official copies of managed objects and data to MOCE 228 for service development purposes; requesting service packages, SIBBs, SLPs or other service or data components 503 from MOCE 228; receiving completed and tested service packages, SIBBs, SLPs or other service or data components 506 from MOCE 228; receiving customer order data 502 from order entry and other legacy systems 229 to provision the IDNA/NGIN system for use by customer; providing unique names to each service component; and, deploying managed objects and data to network service processors via Data Management functions 600, as will be described in greater detail herein.--

---

Please replace the paragraph beginning at page 18, line 31, with the following:

---

A<sup>3</sup> --Other responsibilities of Service Administration include: activating data and service components 512 to ensure that all data, SIBBS and managed objects or service logic programs SLPs are available for nodes via the Data Management component 600; registering the names of the data, SLPs and SIBBs 515 by feeding their logical names to a Network Operating System ("NOS") component 700, to be described in detail below, for registration therewith; distributing data and service components 509; deactivating data and service components 518; and, removing data and services 521 from the IDNA/NGIN system via the Data Management component 600. Service Administration additionally performs a configuration management function by maintaining the state of each SIBB and service (pre-tested, post-tested, deployed, etc.), in

A<sup>3</sup> addition to versioning through its naming process. This ensures a service is not deployed until all components of that service have been successfully tested and configured.--

---

Please replace the paragraph beginning at page 20, line 7, with the following:

---

A<sup>4</sup> --Figure 3(b) illustrates a preferred physical architecture for Service Administration component 500. While Service Administration is a centralized function, it may be embodied as two or more redundant Service Administration sites, e.g., sites 550a, 550b, for reliability with each SA site comprising: SA Servers 560, which may comprise dual redundant processors with a shared disk array 25 comprising the global DBOR 230; and, a personal computer (PC) or workstation 556a,b resident at each respective site 550a, 550b having an interface to enable user access to all Service Administration functions and particularly initiate data and service distribution to specified IDNA/NGIN service nodes, depicted in Figure 3(b) as services nodes 204. The aforementioned data and service distribution activation functions all execute on one or more SA Servers 560 found at each site. The components at each respect SA site 550a,b are connected by an Ethernet LAN 559, which, in turn, is linked to a WAN 566 for communication with the service nodes.--

---

Please replace the paragraph beginning at page 23, line 27, with the following:

---

A<sup>5</sup> --As further shown in Figure 3(c), the SA component 500 comprises the following sub-components: an Inventory Manager 516; a DBOR Manager 520; an Environment Manager 530; an Audit and Reconciliation Manager 535, and, a Monitoring and Logging Manager 540. The functions of each of these will now be explained in greater detail.--

---

Please replace the paragraph beginning at page 26, line 20, with the following:

---

A6 --The Audit and Reconciliation (A/R) Manager 535 ensures data synchronization among the DBOR and its multiple extracts by running auditing routines to compare the data in the DBOR 230 with data in any of various DBOR extracts. It then determines corrective actions to re-sync the multiple databases. To implement these actions, the A/R Manager generates a data package containing data and commands to process this data. This data package is then provided to whichever databases are needed to implement the corrective action to re-sync the multiple databases. Preferably, this may be accomplished as follows: 1) during system idle time, it may run an auditing routine to look for and resolve any discrepancies between the data in the DBOR and the data in a DBOR extract, which may reside in a local Data Management database at a service node; and, 2) during real-time call processing, if a service application finds a discrepancy, e.g., a service application is given a key for a data lookup in Data Management, queries a database with this key, but finds no record, the application generates an alarm. This alarm is sent to the A/R Manager 535, which resolves the discrepancy.--

---

Please replace the paragraph beginning at page 28, line 11, with the following:

---

A7 --Having described the preferred embodiment of the SA component 500, a more detailed description of the major services performed by Service Administration 500, is now provided with reference to Figures 3(c)-3(e).--

---

Please replace the paragraph beginning at page 29, line 9, with the following:

A<sup>8</sup> --As a second major service, service administration component 500 is responsible for service provisioning, i.e., provisioning services with data needed to provide those services. This type of data is input to SA from the Order entry feed 502 and is stored in the global DBOR 230 prior to distribution to Data Management 600. This type of data may include, but is not limited to, customer profile data, such as customer service options, customer name and account data, terminating telephone numbers, call routing data, and any data potentially needed to process and complete a call for a service. As an example, when a 1-800 service is built in Service Creation for a corporate customer, that customer's name, account/billing information, 800 telephone number(s), terminating network addresses, service options (routing features, multi-media file identifiers) received from the OE system are need to provision the particular service(s). In this function, Service Administration 500 parses appropriate order entry feeds to create a consolidated and consistent order entry record to the NGIN and ensures that each feed received from an order entry system or from a provisioning system is acknowledged.--

Please replace the paragraph beginning at page 31, line 11, with the following:

A<sup>9</sup> --Additionally generated in the SA for each service is a service profile, which may be embodied as a formatted data file in SA, that specifies the service's requirements and to which SLEE(s) and/or computers within the network it should be deployed. An example service profile for a particular service to be deployed in the network is depicted in Table 2 as follows:--

Please replace the paragraph beginning at page 32, line 3, with the following:

Ad --In Table 2, there is specified: a service profile name, e.g., service #1001 for a customer X; amount of processing units, memory, and disk space required to execute the service when instantiated; a node instantiate field(s) specifying a time range when a particular service (embodied as a service logic program, for example) is to be instantiated according to a predetermined business rule(s) specified in Service Administration, and a corresponding min/max field(s) indicating the minimum and maximum number of those service objects (SLPs) that may be instantiated by NOS during the specified time range; a special requirements field(s) indicating for example, that the service requires a particular service node capability, e.g., voice playback; and, service start data and service end data. It is readily apparent that SA may distribute the service (and service profile) of the example service 1001 of Table 2 to the service node having the service node profile depicted in Table 1, as the node clearly has the memory requirements and the voice playback support. It is additionally apparent that the example service #1001 depicted in the service profile in Table 2, requires a data set from customer X that would comprise, inter alia, a voice playback service announcement specific to that service #1001 provided by customer X. The SA component 500 will receive data via order entry feed 502 that includes the customer X voice playback announcement, and SA's inventory manager will assign it as a data set #1001, for example, for storage in the DBOR 230. In this manner, SA may automatically distribute the dataset #1001 to the service node(s) providing the service #1001 for customer X.--

Please replace the paragraph beginning at page 39, line 9, with the following:

A<sup>11</sup>  
--According to this fifth SA function, an explanation of how the IDNA/NGIN system handles service construction and deployment phases, is now provided with reference to Figures 3(g) and 3(h) which illustrate a scenario of steps in constructing and deploying an SLP for the IDNA/NGIN system, e.g., for a 1-800 Collect ("1-800-C") service. As indicated at step 812 in Figure 3(g), the MOCE/SCE application program enables the user to access from SA all of the SIBB, SLP, data and other building blocks that are necessary for the creation of the 1-800-C SLP. In the example context of 1-800-C service, such building blocks may include: a play audio building block, a collect digits building block and a voice recognition building block. Copies of these appropriate building blocks are pulled from the global DBOR 230 buy SA into the MOCE/SCE to provide the foundation for developing the 1-800-C Service Logic Program, as indicated at step 814, Figure 3(g). Then, as indicated at step 819, the 1-800-C Service Logic Program and all associated data such as voice files are unit tested within the MOCE/SCE environment. Next, as indicated at step 820, the 1-800-C Service Logic Program will execute correctly once distributed in the network. Then, as indicated at step 823, the 1-800-C Service Logic Program is submitted to the Service whether the 1-800-C SLP can be activated at all repositories where the distribution was successfully received.--

Please replace the paragraph beginning at page 40, line 3, with the following:

A<sup>12</sup>  
--As described herein, the Service Administration component allows the introduction of rules governing data and information distribution, data activation and data removal. Thus, as indicated at step 826, the SA component checks the rules that specify the Data Management repositories that are to receive the Service Logic Program and, the rules regarding the minimum

A<sup>12</sup>  
number of repositories that must receive the distribution prior to allowing activation of the 1-800-C SLP. To do this, as indicated at step 830, Service Administration checks the status of the Data Management repositories by accessing the NOS Network Resource Management function, as described generally herein and in greater detail in U.S. Patent No. 6,425,005 entitled METHOD AND APPARATUS FOR MANAGING RESOURCES IN AN INTELLIGENT NETWORK. Then, as shown at step 832, Figure 3(h), the Service Administration component determines those DM repositories indicating "On-line" status, and, at step 835, distributes the 1-800-C SLP to all the DM repositories that are on-line. For those repositories reporting an off-line status, Service Administration stores the distribution for future forwarding to the off-line repository, as indicated at step 837. Then, as indicated at step 840, the Service Administration component waits until Data Management returns a status for each repository indicating the success or failure of the distribution. A determination is made at step 842 to determine whether the confirmation has been received from the respective DM repository. If the confirmation is not received, the SA waits for the confirmation as indicated at step 844. Once the confirmation is received, the process continues to step 845 where a determination is made by Service Administration as to whether the 1-800-C SLP can be activated at all repositories where the distribution was successfully received.--

---

IN THE CLAIMS:

Please amend claims 1-4 as follows:

---

- A<sup>13</sup>
1. (Amended) A service administration system for distributing service processing resources among one or more service nodes of an intelligent communications network, each